Info-Greedy Sequential Adaptive Compressed Sensing

Gábor Braun, Sebastian Pokutta, Yao Xie

arXiv:1407.0731v4 [cs.IT] 2 Feb 2015

***Abstract*—We present an information-theoretic framework for sequential adaptive compressed sensing, *Info-Greedy Sensing*, where measurements are chosen to maximize the extracted information conditioned on the previous measurements. We show that the widely used bisection approach is Info-Greedy for a family of k-sparse signals by connecting compressed sensing and blackbox complexity of sequential query algorithms, and present Info-Greedy algorithms for Gaussian and Gaussian Mixture Model (GMM) signals, as well as ways to design sparse Info-Greedy measurements. Numerical examples demonstrate the good performance of the proposed algorithms using simulated and real data: Info-Greedy Sensing shows significant improvement over random projection for signals with sparse and low-rank covariance matrices, and adaptivity brings robustness when there is a mismatch between the assumed and the true distributions.**

I. Introduction

Nowadays ubiquitous big data applications (image processing [1], power network monitoring [2], and large scale sensor networks [3]) call for more efficient information *sensing* techniques. Often these techniques are *sequential* in that the measurements are taken one after another. Hence information gained in the past can be used to guide an *adaptive* design of subsequent measurements, which naturally leads to the notion of sequential adaptive sensing. At the same time, a path to efficient sensing of big data is *compressive sensing* [4]–[6], which exploits low-dimensional structures to recover signals from a number of measurements much smaller than the ambient dimension of the signals.

Early compressed sensing works mainly focus on non-adaptive and one-shot measurement schemes. Recently there has also been much interest in *sequential adaptive compressed sensing*, which measures noisy linear combinations of the entries (this is different from the direct adaptive sensing, which measures signal entries directly [7]–[10]). Although in the seminal work of [11], it was shown under fairly general assumptions that "adaptivity does not help much", i.e., sequential adaptive compressed sensing does not improve the *order* of the min-max bounds obtained by algorithms, these limitations are restricted to certain performance metrics. It has also been recognized (see, e.g., [12]–[14]) that adaptive compressed sensing offers several benefits with respect to other performance metrics, such as the reduction in the signal-to-noise ratio (SNR) to recover the signal. Moreover, larger performance gain can be achieved by adaptive compressed sensing if we aim at recovering a "family" of signals with known statistical *prior information* (incorporating statistical priors in compressed sensing has been considered in [15] for the non-sequential setting and in [16] for the sequential setting using Bayesian methods).

To harvest the benefits of adaptive compressed sensing, various algorithms have been developed: compressive binary search [17], [18], which considers a problem of determining the location of a single non-zero entry; a variant of the iterative bisection algorithm [19] to adaptively identify the partial support of the signal; random choice of compressed sensing vectors [20], and a collection of independent structured random sensing matrices in each measurement step [21] with some columns "masked" to zero; an experimental design approach [22] that designs measurements adaptive to the mean square error of the estimated signal; exploiting additional graphical structure of the signal [23], [24]; the CASS algorithm [13], which is based on bisection search to locate multiple non-zero entries, and is claimed to be near-optimal in the number of measurements needed sequentially to achieve small recovery errors; an adaptive sensing strategy specifically tailored to tree-sparse signals [25] that significantly outperforms non-adaptive sensing strategies. In optics literature, compressive imaging systems with sequential measurement architectures have been developed [26]–[28], which may modify the measurement basis based on specific object information derived from the previous measurements and achieve better performance. In medical imaging literature, [29] uses Bayesian experimental design to optimize k-space sampling for nonlinear sparse MRI reconstruction.

The idea of using an information measure for sequential compressed sensing has been spelled out in various places for specific settings or signal models, for example, the seminal Bayesian compressive sensing work [16], which designs a new projection that minimizes the differential entropy of the posterior estimate on a Gaussian signal; [6, Chapter 6.2] and [30], which introduces the so-called "expected information" and outlines a general strategy for sequential adaptive sensing; [31], which develops a two-step adaptive statistical compressed sensing scheme for Gaussian mixture model (GMM) signals based on maximizing an information-theoretic objective function; [32], which sequentially senses low-rank GMM signals based on a posterior distribution and provides an empirical performance analysis; [33] studies the design of linear projection measurements for a vector Poisson signal model; [34] designs general nonlinear functions for mapping high-dimensional data into lower-dimensional space using mutual information as a metric. A general belief, though,

Gábor Braun (Email: gabor.braun@isye.gatech.edu), Sebastian Pokutta (Email: sebastian.pokutta@isye.gatech.edu) and Yao Xie (Email: yao.xie@isye.gatech.edu) are with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

This work is partially supported by NSF grant CMMI-1300144 and CCF-1442635. Authors contributed equally to the paper.

is that it is difficult to devise quantitative error bounds for such sequential information maximizing algorithms (see, e.g., [6, Section 6.2.3]).

In this work, we present a unified information theoretical framework for sequential adaptive compressive sensing, called *Info-Greedy Sensing*, which greedily picks the measurement with the largest amount of information gain based on the previous measurements. More precisely, we design the next measurement to maximize the conditional mutual information between the measurement and the signal with respect to the previous measurements. This framework enables us to better understand existing algorithms, establish theoretical performance guarantees, as well as develop new algorithms. The optimization problem associated with Info-Greedy Sensing is often non-convex. In some cases the solutions can be found analytically, and in others we resort to iterative heuristics. In particular, (1) we show that the widely used bisection approach is Info-Greedy for a family of k-sparse signals by connecting compressed sensing and blackbox complexity of sequential query algorithms [35]; (2) we present Info-Greedy algorithms for Gaussian and Gaussian Mixture Model (GMM) signals under more general noise models (e.g. "noise-folding" [36]) than those considered in [32], and analyze their performance in terms of the number of measurements needed; (3) we also develop new sensing algorithms, e.g., for sparse sensing vectors. Numerical examples are provided to demonstrate the accuracy of theoretical bounds and good performance of Info-Greedy Sensing algorithms using simulated and real data.

The rest of the paper is organized as follows. Section II sets up the formalism for Info-Greedy Sensing. Section III and Section IV present the Info-Greedy Sensing algorithms for k-sparse signals and Gaussian signals (low-rank single Gaussian and GMM), respectively. Section V discusses the Info-Greedy Sensing with sparse measurement vectors. Section VI contains numerical examples using simulated and real data. Finally, Section VII concludes the paper. All proofs are delegated to the Appendix.

The notation in this paper is standard. In particular, $\mathcal{N}(\mu, \Sigma)$ denotes the Gaussian distribution with mean μ and covariance matrix Σ; $[x]_i$ denotes the ith coordinate of the vector x; we use the shorthand $[n] = \{1, \ldots, n\}$; let $|S|$ denote the cardinality of a set S; $\|x\|_0$ is the number of non-zeros in vector x; let $\|\Sigma\|$ be the spectral norm (largest eigenvalue) of a positive definite matrix Σ; let $\det(X)$ be the determinant of a matrix X; let $\mathbb{H}[x]$ denote the entropy of a random variable x; let $\mathbb{I}[x; y]$ denote the mutual information between two random variables x and y. Let the column vector e_i has 1 on the ith entry and zero elsewhere, and let χ^2_n be the quantile function of the chi-squared distribution with n degrees of freedom.

II. Formulation

A typical compressed sensing setup is as follows. Let $x \in \mathbb{R}^n$ be the unknown n-dimensional signal. There are m measurements, and $y \in \mathbb{R}^m$ is the measurement vector depending linearly on the signal x and subject to an additive noise:

$$y = Ax + w, \quad A \triangleq \begin{bmatrix} a_1^\mathsf{T} \\ \vdots \\ a_m^\mathsf{T} \end{bmatrix} \in \mathbb{R}^{m \times n}, \quad w \triangleq \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix} \in \mathbb{R}^{m \times 1}, \tag{1}$$

where $A \in \mathbb{R}^{m \times n}$ is the sensing matrix, and $w \in \mathbb{R}^m$ is the noise vector. Here, each coordinate y_i of y is a result of measuring $a_i^\mathsf{T} x$ with an additive noise w_i. In the setting of sequential compressed sensing, the unknown signal x is measured *sequentially*

$$y_i = a_i^\mathsf{T} x + w_i, \quad i = 1, \ldots, m.$$

In high-dimensional problems, various low-dimensional signal models for x are in common use: (1) sparse signal models, the canonical one being x having $k \ll n$ non-zero entries[1]; (2) low-rank Gaussian model (signal in a subspace plus Gaussian noise); and (3) Gaussian mixture model (GMM) (a model for signal lying in a union of multiple subspaces plus Gaussian noise), which has been widely used in image and video analysis among others[2].

Compressed sensing exploits the low dimensional structure of the signal to recover the signal with high accuracy using much fewer measurements than the dimension of the signal, i.e., $m \ll n$. Two central and interrelated problems in compressed sensing include signal recovery and designing the sensing matrix A. Early compressed sensing works usually assume A to be random, which does have benefits for universality regardless of the signal distribution. However, when there is prior knowledge about the signal distribution, one can optimize A to minimize the number m of measurements subject to a total sensing power constraint

$$\sum_{i=1}^{m} \|a_i\|_2^2 \leq P \tag{2}$$

for some constant $P > 0$. In the following, we either vary power for each measurement $\|a_i\|_2^2 = \beta_i$, or fix them to be unit power $\|a_i\|_2 = 1$ (for example, due to physical constraint) and use repeated measurements β_i times in the direction of a_i, which is equivalent to measuring using an integer valued power. Here β_i can be viewed as the amount of resource we allocated for that measurement (or direction).

We will consider a methodology where A is chosen to extract the most information about the signal, i.e., to maximize *mutual information*. In the non-sequential setting this means that A maximizes the mutual information between the signal x and the measurement outcome, i.e., $A^* = \operatorname{argmax}_A \mathbb{I}[x; Ax + w]$. In sequential compressed sensing, the subsequent measurement vectors can be designed using the already acquired measurements, and hence the sensing matrix A can be designed row by row. Optimal sequential design of A can be defined recursively and viewed as dynamic programming [38]. However, this

[1] In a related model the signal x come from a dictionary with few nonzero coefficients, whose support is unknown. We will not further consider this model here.

[2] A mixture of GMM models has also been used to study sparse signals [37]. There are also other low-dimensional signal models including the general manifold models which will not be considered here.

formulation is usually intractable in all but the most simple situations (one such example is the sequential probabilistic bisection algorithm in [30], which locates a single non-zero entry). Instead, the usual approach operates in a greedy fashion. The core idea is that based on the information that the previous measurements have extracted, the new measurement should probe in the direction that maximizes the conditional information as much as possible. We formalize this idea as *Info-Greedy Sensing*, which is described in Algorithm 1. The algorithm is initialized with a prior distribution of signal x, and returns the Bayesian posterior mean as an estimator for signal x. Conditional mutual information is a natural metric, as it counts only useful new information between the signal and the potential result of the measurement disregarding noise and what has already been learned from previous measurements.

Algorithm 1 Info-Greedy Sensing

Require: distributions of signal x and noise w, error tolerance ε or maximum number of iterations M

1: $i \leftarrow 1$
2: **repeat**
3: $\quad a_i \leftarrow \operatorname{argmax}_{a_i} \mathbb{I}\left[x; a_i^\intercal x + w_i \,|\, y_j, a_j, j < i\right] / \beta_i$
4: $\quad y_i = a_i^\intercal x + w_i$ {measurement}
5: $\quad i \leftarrow i + 1$
6: **until** $\mathbb{I}\left[x; a_i^\intercal x + w_i \,|\, y_j, a_j, j \le i\right] \le \delta(\varepsilon)$ or $i > M$.

Algorithm 1 stops either when the conditional mutual information is smaller than a threshold $\delta(\varepsilon)$, or we have reached the maximum number of iterations M. How $\delta(\varepsilon)$ relates to the precision ε depends on the specific signal model employed. For example, for Gaussian signal, the conditional mutual information is the log determinant of the conditional covariance matrix, and hence the signal is constrained to be in a small ellipsoid with high probability. Also note that in this algorithm, the recovered signal may not reach accuracy ε if it exhausts the number of iterations M. In theoretical analysis we assume M is sufficiently large to avoid it.

Note that the optimization problem in Info-Greedy Sensing $\operatorname{argmax}_{a_i} \mathbb{I}\left[x; a_i^\intercal x + w_i \,|\, y_j, a_j, j < i\right]$ is non-convex in general [39]. Hence, we will discuss various heuristics and establish their theoretical performance in terms of the following metric:

Definition II.1 (Info-Greedy)**.** We call an algorithm *Info-Greedy* if the measurement maximizes $\mathbb{I}\left[x; y_i \,|\, y_j : j < i\right] / \beta_i$ for each i, where x is the unknown signal, y_j is the measurement outcome, and β_i is the amount of resource for measurement i.

III. k-SPARSE SIGNAL

In this section, we consider the Info-Greedy Sensing for k-sparse signal with arbitrary nonnegative amplitudes in the noiseless case as well as under Gaussian measurement noise. We show that a natural modification of the bisection algorithm corresponds to Info-Greedy Sensing under a certain probabilistic model. We also show that Algorithm 2 is optimal in terms of the number of measurements for 1-sparse signals as well as optimal up to a $\log k$ factor for k-sparse signals in the noiseless case. In the presence of Gaussian measurement noise, it is optimal up to at most another $\log n$ factor. Finally, we show Algorithm 2 is Info-Greedy when $k = 1$, and when $k > 1$ it is Info-Greedy up to a $\log k$ factor.

To simplify the problem, we assume the sensing matrix A consists of binary entries: $a_{ij} \in \{0, 1\}$. Consider a signal with each element $x_i \in \mathbb{R}_+$ with up to k non-zero entries which are distributed uniformly at random. The following lemma gives an upper bound on the number of measurements m for our modified bisection algorithm (see Algorithm 2) to recover such x. In the description of Algorithm 2, let

$$[a_S]_i \coloneqq \begin{cases} 1, & i \in S \\ 0, & i \notin S \end{cases}$$

denote the characteristic vector of a set S. The basic idea is to recursively estimate a tuple (S, ℓ) that consists of a set S which contains possible locations of the non-zero elements, and the total signal amplitude in that set. We say that a signal x has minimum amplitude τ, if $x_i > 0$ implies $x_i \ge \tau$ for all $i \in [n]$.

Theorem III.1 (Upper bound for k-sparse signal x)**.** *Let $x \in \mathbb{R}_+^n$ be a k-sparse signal.*

1) *In the noiseless case, Algorithm 2 recovers the signal x exactly with at most $2k\lceil \log n \rceil$ measurements (using $r = 1$ in Line 1).*
2) *In the noisy case with $w_i \sim \mathcal{N}(0, \sigma^2)$, Algorithm 2 recovers the signal x such that $\|x - \widehat{x}\|_2 \le \sqrt{k}\varepsilon$ with probability at least $1 - k\lceil \log n \rceil / (n^{\varepsilon^2/(2k\sigma^2)}) = O(1)$ using at most $2k\lceil \log n \rceil^2$ measurements.*

Algorithm 2 Bisection for k-sparse signals

Require: ambient dimension n of x, error probability δ, noise variance σ, error ε

1: $r \leftarrow \lceil \log n \rceil$
2: $L \leftarrow \{[n]\}$
3: $\widehat{x} \leftarrow 0$ {initialize estimator}
4: **while** L not empty **do**
5: $\quad$ **for all** $S \in L$ **do**
6: $\quad\quad$ Partition $S = S_1 \dot\cup S_2$ with $||S_1| - |S_2|| \le 1$
7: $\quad\quad$ Replace S by S_1 and S_2 in L
8: $\quad$ **end for**
9: $\quad$ **for all** $S \in L$ **do**
10: $\quad\quad$ Measure r times and average: $y = a_S^\intercal x + w$
11: $\quad\quad$ **if** $y \le \varepsilon$ **then**
12: $\quad\quad\quad$ Remove S from L. {$\widehat{x}$ is already 0 on S.}
13: $\quad\quad$ **else if** $|S| = 1$ **then**
14: $\quad\quad\quad$ Remove S from L.
15: $\quad\quad\quad$ $\widehat{x}_i \leftarrow y$ where $S = \{i\}$.
16: $\quad\quad$ **end if**
17: $\quad$ **end for**
18: **end while**
19: **return** $\widehat{x}$ as estimator for x.

Lemma III.2 (Lower bound for noiseless k-sparse signal x)**.** *Let $x \in \mathbb{R}_+^n$, $x_i \in \{0, 1\}$ be a k-sparse signal. Then to recover*

x exactly, the expected number of measurements m required for any algorithm is at least $\frac{k}{\log k+1}(-1+\log n)$.

Lemma III.3 (Bisection Algorithm 2 for $k=1$ is Info-Greedy)**.** *For $k=1$ Algorithm 2, is Info-Greedy.*

In general case the simple analysis that leads to Lemma III.3 fails. However, using Theorem A.1 in the Appendix we can estimate the average amount of information obtained from a measurement:

Lemma III.4 (Bisection Algorithm 2 is Info-Greedy up to a $\log k$ factor in the noiseless case)**.** *Let $k \leq n \in \mathbb{N}$. Then the average information of a measurement in Algorithm 2:*

$$\mathbb{I}\left[X; Y_i \,|\, Y_1, \ldots, Y_{i-1}\right] \geq 1 - \frac{\log k}{\log n}.$$

Remark III.5. 1) Observe that Lemma III.4 establishes that Algorithm 2 for a sparse signal with $\log k = o(\log n)$ acquires at least a $\frac{1}{\log k+1} - o(1)$ fraction of the maximum possible mutual information (which on average is roughly 1 bit per measurement).

2) Here we constrained the entries of matrix A to be binary valued. This may correspond to applications, for examples, sensors reporting errors and the measurements count the total number of errors. Note that, however, if we relax this constraint and allow entries of A to be real-valued, in the absence of noise the signal can be recovered from one measurement that project the signal onto a vector with entries $[2^0, 2^1, 2^2, \cdots]$.

3) The setup here with k-sparse signals and binary measurement matrix A generalizes the group testing [40] setup.

4) the CASS algorithm [13] is another algorithm that recovers a k-sparse signal x by iteratively partitioning the signal support into $2k$ subsets, computing the sum over that subset and keeping the largest k. In [13] it was shown that to recover a k-sparse x with non-uniform positive amplitude with high probability, the number of measurements m is on the order of $2k\log(n/k)$ with *varying* power measurement. It is important to note that the CASS algorithm allows for power allocation to mitigate noise, while we repeat measurements. This, however, coincides with the number of unit length measurements of our algorithm, $2k\lceil\log n\rceil^2$ in Lemma III.1 after appropriate normalization. For specific regimes of error probability, the $O(\log n)$ overhead in Lemma III.1 can be further reduced. For example, for any constant probability of error $\varepsilon > 0$, the number of required repetitions per measurement is $O(\log\log n)$ leading to improved performance. Our algorithm can be also easily modified to incorporate power allocation.

IV. Low-Rank Gaussian Models

In this section, we derive the Info-Greedy Sensing algorithms for the single low-rank Gaussian model as well as the low-rank GMM signal model, and also quantify the algorithm's performance.

A. Single Gaussian model

Consider a Gaussian signal $x \sim \mathcal{N}(\mu, \Sigma)$ with known parameters μ and Σ. The covariance matrix Σ has rank $k \leq n$. We will consider three noise models:

1) white Gaussian noise added after the measurement (the most common model in compressed sensing):
$$y = Ax + w, \quad w \sim \mathcal{N}(0, \sigma^2 I). \tag{3}$$
Let $\beta_i = \|a_i\|_2^2$ represent the power allocated to the ith measurement. In this case, higher power β_i allocated to a measurement increases SNR of that measurement.

2) white Gaussian noise added prior to the measurement, a model that appears in some applications such as reduced dimension multi-user detection in communication systems [41] and also known as the "noise folding" model [36]:
$$y = A(x + w), w \sim \mathcal{N}(0, \sigma^2 I). \tag{4}$$
In this case, allocating higher power for a measurement cannot increase the SNR of the outcome. Hence, we use the actual number of repeated measurements in the same direction as a proxy for the amount of resource allocated for that direction.

3) colored Gaussian noise with covariance Σ_w added either prior to the measurement:
$$y = A(x + w), w \sim \mathcal{N}(0, \Sigma_w), \tag{5}$$
or after the measurement:
$$y = Ax + w, w \sim \mathcal{N}(0, \Sigma_w). \tag{6}$$

In the following, we will establish lower bounds on the amount of resource (either the minimum power or the number of measurements) needed for Info-Greedy Sensing to achieve a recovery error $\|x - \widehat{x}\|_2 \leq \varepsilon$.

1) White noise added prior to measurement or "noise folding": We start our discussion with this model and results for other models can be derived similarly. As β_i does not affect SNR, we set $\|a_i\|_2 = 1$. Note that conditional distribution of x given y_1 is a Gaussian random vector with adjusted parameters

$$\begin{aligned} x \mid y_1 &\sim \mathcal{N}(\mu + \Sigma a_1 (a_1^\intercal \Sigma a_1 + \sigma^2)^{-1}(y_1 - a_1^\intercal \mu), \\ &\quad \Sigma - \Sigma a_1 (a_1^\intercal \Sigma a_1 + \sigma^2)^{-1} a_1^\intercal \Sigma). \end{aligned} \tag{7}$$

Therefore, to find Info-Greedy Sensing for a single Gaussian signal, it suffices to characterize the first measurement $a_1 = \arg\max_{a_1} \mathbb{I}\left[x; y_1\right]$ and from there on iterate with adjusted distributional parameters. For Gaussian signal $x \sim \mathcal{N}(\mu, \Sigma)$ and the noisy measurement $\sigma > 0$, we have

$$\mathbb{I}\left[x; y_1\right] = \mathbb{H}\left[y_1\right] - \mathbb{H}\left[y_1 \,|\, x\right] = \frac{1}{2}\ln\left(a_1^\intercal \Sigma a_1/\sigma^2 + 1\right). \tag{8}$$

Clearly, with $\|a_1\|_2 = 1$, (8) is maximized when a_1 corresponds to the largest eigenvector of Σ. From the above argument, the Info-Greedy Sensing algorithm for a single Gaussian signal is to choose $a_1, a_2, \ldots$ as the orthonormal eigenvectors of Σ in a decreasing order of eigenvalues, as described in Algorithm 3. The following theorem establishes the bound on the number of measurements needed.

Theorem IV.1 (White Gaussian noise added prior to measurement or "noise folding")**.** *Let $x \sim \mathcal{N}(\mu, \Sigma)$ and let $\lambda_1, \ldots, \lambda_k$ be the eigenvalues of Σ with multiplicities. Further let $\varepsilon > 0$ be the accuracy and $w_i \sim \mathcal{N}(0, \sigma^2)$. Then Algorithm 3 recovers x satisfying $\|x - \widehat{x}\|_2 < \varepsilon$ with probability at least p using at most the following number of measurements by unit vectors $\|a_i\|_2 = 1$:*

$$m = \sum_{\substack{i=1 \\ \lambda_i \neq 0}}^{k} \max\left\{0, \left\lceil \left(\frac{\chi_n^2(p)}{\varepsilon^2} - \frac{1}{\lambda_i} \right) \sigma^2 \right\rceil \right\} \tag{9a}$$

provided $\sigma > 0$. If $\sigma^2 \leq \varepsilon^2 / \chi_n^2(p)$ the number of measurements simplifies to

$$\left| \left\{ i : \lambda_i > \frac{\varepsilon^2}{\chi_n^2(p)} \right\} \right|. \tag{9b}$$

This also holds when $\sigma = 0$.

2) White noise added after measurement: A key insight in the proof for Theorem IV.1 is that repeated measurements in the same eigenvector direction corresponds to a single measurement in that direction with all the power summed together. This can be seen from the following discussion. After measuring in the direction of a unit norm eigenvector u with eigenvalue λ, and using power β, the conditional covariance matrix takes the form of

$$\begin{aligned} &\Sigma - \Sigma\sqrt{\beta}u \left(\sqrt{\beta} u^\intercal \Sigma \sqrt{\beta} u + \sigma^2 \right)^{-1} \sqrt{\beta} u^\intercal \Sigma \\ &= \frac{\lambda \sigma^2}{\beta\lambda + \sigma^2} u u^\intercal + \Sigma^{\perp u}, \end{aligned} \tag{10}$$

where $\Sigma^{\perp u}$ is the component of Σ in the orthogonal complement of u. Thus, the only change in the eigendecomposition of Σ is the update of the eigenvalue of u from λ to $\lambda\sigma^2/(\beta\lambda + \sigma^2)$. Informally, measuring with power allocation β on a Gaussian signal x reduces the uncertainty in direction u as illustrated in Fig. 1. We have the following performance bound for sensing a Gaussian signal:

Theorem IV.2 (White Gaussian noise added after measurement)**.** *Let $x \sim \mathcal{N}(\mu, \Sigma)$ and let $\lambda_1, \ldots, \lambda_k$ be the eigenvalues of Σ with multiplicities. Further let $\varepsilon > 0$ be the accuracy and $w_i \sim \mathcal{N}(0, \sigma^2)$. Then Algorithm 3 recovers x satisfying $\|x - \widehat{x}\|_2 < \varepsilon$ with probability at least p using at most the following power*

$$P = \sum_{\substack{i=1 \\ \lambda_i \neq 0}}^{k} \max\left\{0, \left(\frac{\chi_n^2(p)}{\varepsilon^2} - \frac{1}{\lambda_i} \right) \sigma^2 \right\} \tag{11}$$

provided $\sigma > 0$.

3) Colored noise: When a colored noise $w \sim \mathcal{N}(0, \Sigma_w)$ is added either prior to, or after the measurement, similar to the white noise cases, the conditional distribution of x given the first measurement y_1 is a Gaussian random variable with adjusted parameters. Hence, as before, the measurement vectors can be found iteratively. Algorithm 3 presents Info-Greedy Sensing for this case and the derivation is given in Appendix B. Algorithm 3 also summarizes all the Info-Greedy Sensing algorithms for Gaussian signal under various noise models.

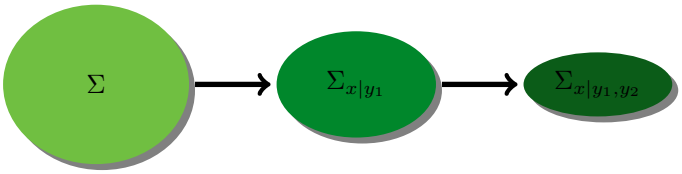

Fig. 1. Evolution of the covariance matrix by sequentially measuring with an eigenvector of the largest eigenvalue.

Algorithm 3 Info-Greedy Sensing for Gaussian signals

Require: signal mean μ and covariance Σ, accuracy ε, probability of correctness p, noise covariance matrix Σ_w (for white noise $\sigma^2 I$)

repeat
  if white noise added after measurement **then**
    $\lambda \leftarrow \|\Sigma\|$ {largest eigenvalue}
    $u \leftarrow$ eigenvector of Σ for eigenvalue λ
    $\beta \leftarrow \left(\frac{\chi_n^2(p)}{\varepsilon^2} - \frac{1}{\lambda} \right) \sigma^2$
    $a \leftarrow \sqrt{\beta} u$
    $y = a^\intercal x + w$
  else if white noise added prior to measurement **then**
    $\lambda \leftarrow \|\Sigma\|$ {largest eigenvalue}
    $u \leftarrow$ eigenvector of Σ for eigenvalue λ
    $a \leftarrow \sqrt{\beta} u$
    $y = a^\intercal (x + w)$
  else if colored noise added after measurement **then**
    $\Sigma = U_x \Lambda_x U_x^\intercal, \Sigma_w = U_w \Lambda_w U_w^\intercal$ {eigendecomposition}
    $u \leftarrow (1/\|\Lambda_w^{1/2} U_w^\intercal e_1\|_2) U_x \Lambda_w^{1/2} U_w^\intercal e_1$
    $a \leftarrow \sqrt{\beta} u$
    $y = a^\intercal x + w$
  else if colored noise added prior to measurement **then**
    $\lambda \leftarrow \|\Sigma_w^{-1} \Sigma\|$ {largest eigenvalue}
    $\beta \leftarrow \frac{\chi_n^2(p)}{\varepsilon^2} \|\Sigma_w\| - \frac{1}{\lambda}$
    $u \leftarrow$ largest eigenvector of $\Sigma_w^{-1} \Sigma$ for eigenvalue λ
    $a \leftarrow \sqrt{\beta} u$
    $y = a^\intercal (x + w)$
  end if
  $\mu \leftarrow \mu + \Sigma a (a^\intercal \Sigma a + \sigma^2)^{-1} (y - a^\intercal \mu)$ {mean}
  $\Sigma \leftarrow \Sigma - \Sigma a (a^\intercal \Sigma a + \sigma^2)^{-1} a^\intercal \Sigma$ {covariance}
until $\|\Sigma\| \leq \varepsilon^2 / \chi_n^2(p)$ {all eigenvalues become small}
return posterior mean μ

The following version of Theorem IV.1 is for the required number of measurements for colored noise in the "noise folding" model:

Theorem IV.3 (Colored Gaussian noise added prior to measurement or "noise folding")**.** *Let $x \sim \mathcal{N}(\mu, \Sigma)$ be a Gaussian signal, and let $\lambda_1, \ldots, \lambda_n$ denote the eigenvalues of $\Sigma_w^{-1}\Sigma$ with multiplicities. Assume $w \sim \mathcal{N}(0, \Sigma_w)$. Furthermore, let $\varepsilon > 0$ be the required accuracy. Then Algorithm 3 recovers x satisfying $\|x - \widehat{x}\|_2 < \varepsilon$ with probability at least p using at most the following number of measurements by unit vectors*

$\|a_i\|_2 = 1$:

$$m = \sum_{\substack{i=1 \\ \lambda_i \neq 0}}^{n} \max\left\{0, \left\lceil \frac{\chi_n^2(p)}{\varepsilon^2}\|\Sigma_w\| - \frac{1}{\lambda_i} \right\rceil \right\}. \quad (12)$$

Remark IV.4. (1) Under these noise models, the posterior distribution of the signal is also Gaussian, and the measurement outcome y_i affects *only its mean* and but not the covariance matrix (see (7)). In other words, the outcome does not affect the mutual information of posterior Gaussian signal. In this sense, for Gaussian signals adaptivity brings no advantage *when* Σ *is accurate*, as the measurements are pre-determined by the eigenspace of Σ. However, when knowledge of Σ is inaccurate for Gaussian signals, adaptivity brings benefit as demonstrated in Section VI-A1, since a sequential update of the covariance matrix incorporates new information and "corrects" the covariance matrix when we design the next measurement.

(2) In (10) the eigenvalue λ reduces to $\lambda\sigma^2/(\beta\lambda + \sigma^2)$ after the first measurement. Now iterating this we see by induction that after m' measurements in direction a, the eigenvalue λ reduces to $\lambda\sigma^2/(m'\beta\lambda + \sigma^2)$, which is the same as measuring once in direction a with power $m'\beta$. Hence, measuring several times in the same direction of a, and thereby splitting power into $\beta_1, \ldots, \beta_{m'}$ for the measurements, has the same effect as making one measurement with total the power $\sum_{i=1}^{m'} \beta_i$.

(3) Info-Greedy Sensing for Gaussian signal can be implemented efficiently. Note that in the algorithm we only need compute the leading eigenvector of the covariance matrix; moreover, updates of the covariance matrix and mean are simple and iterative. In particular, for a sparse $\Sigma \in \mathbb{R}^{n\times n}$ with v non-zero entries, the computation of the largest eigenvalue and associated eigenvector can be implemented in $\mathcal{O}(t(n+v))$ using sparse power's method [42], where t is the number of power iterations. In many high-dimensional applications, Σ is sparse if the variables (entries of x) are not highly correlated. Also note that the sparsity structure of the covariance matrix as well as the correlation structure of the signal entries will not be changed by the update of the covariance matrix. This is because in (10) the update only changes the eigenvalues but not the eigenvectors. To see why this is true, let $\Sigma = \sum_i \lambda_i q_i q_i^\top$ be the eigendecomposition of Σ. By saying that the covariance matrix is sparse, we assume that q_i's are sparse and, hence, the resulting covariance matrix Σ has few non-zero entries. Therefore, updating the covariance matrix will not significantly change the number of non-zero entries in a covariance matrix. We demonstrate the scalability of Info-Greedy Sensing with larger examples in Section VI-A1.

B. Gaussian mixture model (GMM)

The probability density function of GMM is given by

$$p(x) = \sum_{c=1}^{C} \pi_c \mathcal{N}(\mu_c, \Sigma_c), \quad (13)$$

where C is the number of classes, and π_c is the probability of samples from class c. Unlike Gaussian, mutual information of GMM cannot be explicitly written. However, for GMM signals a gradient descent approach that works for an arbitrary signal model can be used as outlined in [32]. The derivation uses the fact that the gradient of the conditional mutual information with respect to a_i is a linear transform of the minimum mean square error (MMSE) matrix [39], [43]. Moreover, the gradient descent approach for GMM signals exhibits structural properties that can be exploited to reduce the computational cost for evaluating the MMSE matrix, as outlined in [32], [37]. For completeness we include the detail of the algorithm here, as summarized in Algorithm 6 and the derivations are given in Appendix C[3].

An alternative heuristic for sensing GMM is the so-called *greedy heuristic*, which is also mentioned in [32]. The heuristic picks the Gaussian component with the highest posterior π_c at that moment, and chooses the next measurement a to be its eigenvector associated with the maximum eigenvalue, as summarized in Algorithm 6. The greedy heuristic is not Info-Greedy, but it can be implemented more efficiently compared to the gradient descent approach. The following theorem establishes a simple upper bound on the number of required measurements to recover a GMM signal using the greedy heuristic with small error. The analysis is based on the well-known multiplicative weight update method (see e.g., [45]) and utilizes a simple reduction argument showing that when the variance of every component has been reduced sufficiently to ensure a low error recovery with probability p, we can learn (a mix of) the right component(s) with few extra measurements.

Theorem IV.5 (Upper bound on m of greedy heuristic algorithm for GMM)**.** *Consider a GMM signal x parameterized in (13). Let m_c be the required number of measurements (or power) to ensure $\|x - \widehat{x}\|_2 < \varepsilon$ with probability p for a Gaussian signal $\mathcal{N}(\mu_c, \Sigma_c)$ corresponding to component c for all $c \in C$. Then we need at most*

$$\left(\sum_{c\in C} m_c\right) + \Theta(\frac{1}{\tilde{\eta}} \ln |C|)$$

measurements (or power) to ensure $\|x - \widehat{x}\|_2 < \varepsilon$ when sampling from the posterior distribution of π with probability $p(1 - \tilde{\eta} - o(1))$.

Remark IV.6. In the high noise case, i.e., when SNR is low, Info-Greedy measurements can be approximated easily. Let c_0 denote the random variable indicating the class where the signal is sampled from. Then $\mathbb{I}[x; y] = \mathbb{I}[x; y \,|\, c] + \mathbb{I}[x; c] - \underbrace{\mathbb{I}[x; c \,|\, y]}_{0} = \mathbb{I}[x; c] + \sum_c \pi_c \log(1 + a^\intercal \Sigma_c a/\sigma^2)/2 \propto \sum_c \pi_c a^\intercal \Sigma_c a/\sigma^2 = a^\intercal(\sum_c \pi_c \Sigma_c)a/\sigma^2$. Hence, the Info-Greedy measurement should be the leading eigenvector of the average covariance matrix with the posterior weights.

V. Sparse measurement vector

In various applications, we are interested in finding a sparse measurement vector a. With such requirement, we can add a cardinality constraint on a in the Info-Greedy Sensing

[3]Another related work is [44] which studies the behavior of minimum mean sure error (MMSE) associated with the reconstruction of a signal drawn from a GMM as a function of the properties of the linear measurement kernel and the Gaussian mixture, i.e. whether the MMSE converges or does not converge to zero as the noise.

Algorithm 4 Gradient descent for mutual information maximizing measurement

Require: initial a_i, step-size μ, tolerance $\eta > 0$

repeat
  generate $c_0 \sim \widetilde{\pi}_c$, and $x_0 \sim \mathcal{N}(\mu_{c_0} + \Sigma_c D_{i-1}(y_{-i} - D_{i-1}\mu_c)/\sigma^2, \Sigma_c - \Sigma_c D_{i-1}^{\mathsf{T}} D_{i-1}\Sigma_c/\sigma^2)$.
  measure $y_0 = a_i^{\mathsf{T}} x_0 + w_i$
  evaluate $g(y_0)$ using (28)
  estimate $E_i \approx \frac{1}{N}\sum_{j=1}^{N} \tilde{p}(y_j) g(y_j)$.
  evaluate $h_i(a_i) \triangleq \partial \mathbb{I}\left[x; y_i \,|\, y_j, j < i\right] / \partial a_i$ using (23)
  update $a_i \leftarrow a_i + \mu h_i(a_i)$
  evaluate approximated mutual information using (29)
until increase in mutual information $\leq \eta$
return solution measurement vector a_i

Algorithm 5 Update GMM distributional parameters

Require: mean $\{\mu_c\}$, covariance $\{\Sigma_c\}$, number of GMM components C, distribution $\{\pi_c\}$, standard deviation σ of noise, matrix contains vectors thus far D and measurements acquired thus far $\widetilde{y}$

for $c = 1, \ldots, C$ **do**
  $\mu_c \leftarrow \mu_c + \Sigma_c a (a^{\mathsf{T}}\Sigma_c a + \sigma^2)^{-1}(y - a^{\mathsf{T}}\mu_c)$ {mean}
  $\Sigma_c \leftarrow \Sigma_c - \Sigma_c a (a^{\mathsf{T}}\Sigma_c a + \sigma^2)^{-1} a^{\mathsf{T}}\Sigma_c$ {covariance}
  $\pi_c \leftarrow \pi_c \mathcal{N}(\widetilde{y}, D\mu_c D\Sigma_c D^{\mathsf{T}} + \sigma^2)$
end for
$\pi_c \leftarrow \pi_c / \sum_{c=1}^{C} \pi_c$ {normalizing distribution}
return updated parameters $\{\mu_c, \Sigma_c, \pi_c\}$

Algorithm 6 Info-Greedy Sensing for GMM using greedy heuristic and gradient descent approach

Require: mean $\{\mu_c\}_{c=1}^{C}$, covariance $\{\Sigma_c\}_{c=1}^{C}$, initial distribution $\{\pi_c\}_{c=1}^{C}$ standard deviation σ of noise, probability of correctness p

Initialize $\mu_c^{(0)} = \mu_c$, $\Sigma_c^{(0)} = \Sigma_c$, $\pi_c^{(0)} = \pi_c$
repeat
  if greedy heuristic **then**
    $z \leftarrow \arg\max_c \pi_c^{(i-1)}$
    $a_i \leftarrow$ largest eigenvector of $\Sigma_z^{(i-1)}$
  else if gradient decent approach **then**
    $a_i \leftarrow$ solved from Algorithm 4
  end if
  $y_i = a_i^{\mathsf{T}} x + w_i$ {measure}
  update parameters $\mu_c^{(i)}, \Sigma_c^{(i)}, \pi_c^{(i)}$ using Algorithm 5
until reach maximum iteration
return signal estimate $c^* = \arg\max_c \pi_c^{(I)}$, $\widehat{\mu} = \mu_{c^*}^{(I)}$

formulation: $\|a\|_0 \leq k_0$, where k_0 is the number of non-zero entries we allowed for a vector. This is a non-convex integer program with non-linear cost function, which can be solved by outer approximation [46], [47]. The idea of outer approximation is to generate a sequence of cutting planes to approximate the cost function via its subgradient and iteratively include these cutting planes as constraints in the original optimization problem. In particular, we initialize by solving the following optimization problem

$$\begin{array}{ll} \underset{a,r,z}{\text{maximize}} & z \\ \text{subject to} & \sum_{i=1}^{n} r_i \leq k_0 \\ & a_i \leq r_i, \quad -a_i \leq r_i \\ & 0 \leq z \leq c, \quad r_i \in \{0,1\}, i = 1, \ldots, n \\ & a \in \mathbb{R}^n, \quad z \in \mathbb{R}, \end{array} \tag{14}$$

where r and z are introduced auxiliary variables, and c is an user specified upper bound that bounds the cost function over the feasible region. The constraint of the above optimization problem can be casted into matrix vector form as follows:

$$F_0 \triangleq \left[\begin{array}{c|c|c} 1_{1\times n} & 0_{1\times n} & 0 \\ \hline -I_n & I_n & 0_{n\times 1} \\ \hline -I_n & -I_n & 0_{n\times 1} \\ \hline 0_{1\times n} & 0_{1\times n} & 1 \\ \hline 0_{1\times n} & 0_{1\times n} & -1 \end{array}\right], \quad g_0 \triangleq \begin{bmatrix} k_0 \\ 0_{2n\times 1} \\ c \\ 0 \end{bmatrix}$$

such that $F_0 \begin{bmatrix} r & a & z \end{bmatrix}^{\mathsf{T}} \leq g_0$. The mixed-integer linear program formulated in (14) can be solved efficiently by a standard software such as GUROBI[4]. In the next iteration, solution a_* to this optimization problem will be used to generate a new cutting plane, which we include in the original problem by appending a row to F_ℓ and adding an entry to g_ℓ as follows

$$F_{\ell+1} = \left[\begin{array}{c} F_\ell \\ \hline 0 \quad -(\nabla f(a_*))^{\mathsf{T}} \quad 1 \end{array}\right], \tag{15}$$

$$g_{\ell+1} = \left[\begin{array}{c} g_\ell \\ \hline f(a_*) - a_*^{\mathsf{T}} \nabla f(a_*) \end{array}\right], \tag{16}$$

where f is the non-linear cost function in the original problem. For Gaussian signal x, the cost function and its gradient take the form of:

$$f(a) = \frac{1}{2}\log(\frac{a^{\mathsf{T}}\Sigma a}{\sigma^2} + 1), \quad \nabla f(a) = \frac{1}{a^{\mathsf{T}}\Sigma a + \sigma^2}\Sigma a. \tag{17}$$

By repeating iterations as above, we can find a measurement vector with sparsity k_0 which is approximately Info-Greedy.

VI. Numerical examples

A. Simulated examples

1) Low-rank Gaussian model: First, we examine the performance of Info-Greedy Sensing for Gaussian signal. The dimension of the signal is $n = 100$, and we set the probability of recovery $p = 0.95$, the noise standard deviation $\sigma = 0.01$. The signal mean vector $\mu = 0$, where the covariance matrix Σ is generated as $\Sigma = \mathcal{T}_{0.7}(\Sigma_0\Sigma_0^{\mathsf{T}}/\|\Sigma_0\Sigma_0^{\mathsf{T}}\|_2)$, $\Sigma_0 \in \mathbb{R}^{n\times n}$ has each entry i.i.d. $\mathcal{N}(0,1)$, and the operator $\mathcal{T}_{0.7}(X)$ thresholds eigenvalues of a matrix X that are smaller than 0.7 to be zero.

[4] http://www.gurobi.com

The error tolerance $\epsilon = 0.1$ (represented as dashed lines in the figures). For the white noise case, we set $w \sim \mathcal{N}(0, \sigma^2 I)$, and for the colored noise case, $w \sim \mathcal{N}(0, \Sigma_w)$ and the noise covariance matrix Σ_w is generated randomly as $\tilde{\Sigma}_0^\intercal \tilde{\Sigma}_0 / \|\tilde{\Sigma}_0^\intercal \tilde{\Sigma}_0\|_2$ for a random matrix $\tilde{\Sigma}_0$ with entries i.i.d. $\mathcal{N}(0,1)$. The number of measurements is determined from Theorem IV.1 and Theorem IV.2. We run the algorithm over 1000 random instances. Fig. 2 demonstrates the ordered recovery error $\|x - \widehat{x}\|_2$, as well as the ordered number of measurements calculated from the formulas, for the white and colored noise cases, respectively. Note that in both the white noise and colored noise cases, the errors for Info-Greedy Sensing can be two orders of magnitude lower than the errors obtained from measurement using Gaussian random vectors, and the errors fall below our desired tolerance ε using the theoretically calculated m.

When the assumed covariance matrix for the signal x is equal to its true covariance matrix, Info-Greedy Sensing is identical to the batch method [32] (the batch method measures using the largest eigenvectors of the signal covariance matrix). However, when there is a mismatch between the two, Info-Greedy Sensing outperforms the batch method due to adaptivity, as shown in Fig. 3. For Gaussian signals, the complexity of the batch method is $\mathcal{O}(n^3)$ (due to eigendecomposition), versus the complexity of Info-Greedy Sensing algorithm is on the order of $\mathcal{O}(tmn^2)$ where t is the number of iterations needed to compute the eigenvector associated with the largest eigenvalue (e.g., using the power method), and m is the number of measures which is typically on the order of k.

We also try larger examples. Fig. 4 demonstrates the performance of Info-Greedy Sensing for a signal x of dimension 1000 and with dense and low-rank Σ (approximately 5% of non-zero eigenvalues). Another interesting case is shown in Fig. 5, where $n = 5000$ and Σ is rank 3 and very sparse: only about 0.0003% of the entries of Σ are non-zeros. In this case Info-Greedy Sensing is able to recover the signal with a high precision using only 3 measurements. This shows the potential value of Info-Greedy Sensing for big data.

2) Low-rank GMM model: In this example we consider a GMM model with $C = 3$ components, and each Gaussian component is generated as a single Gaussian component described in the previous example Section VI-A1 ($n = 100$ and $\sigma = 0.01$). The true prior distribution is $\pi = (0.3, 0.2, 0.5)$ for the three components (hence each time the signal x is draw from one component with these probabilities), and the assumed prior distribution for the algorithms is uniform: each component has probability $1/3$. The parameters for the gradient descent approach are: step size $\mu = 0.2$ and the error tolerance to stop the iteration $\eta = 0.01$. Fig. 6 demonstrates the estimated cumulative mutual information and mutual information in a single step, averaged over 100 Monte Carlo trials, and the gradient descent based approach has higher information gain than that of the greedy heuristic, as expected. Fig. 7 shows the ordered errors for the batch method based on mutual information gradient [32], the greedy heuristic versus gradient descent approach, when $m = 11$ and $m = 20$, respectively. Note that Info-Greedy Sensing approaches (greedy heuristic and gradient descent) outperform the batch method due to adaptivity, and that the simpler greedy heuristic performs fairly

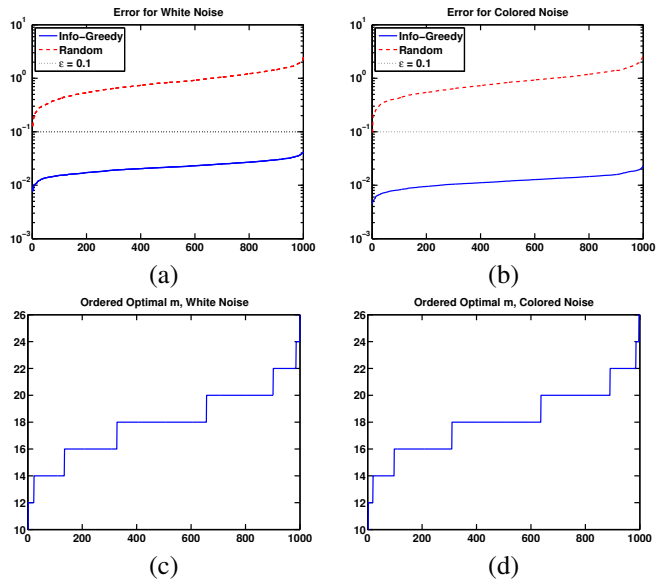

Fig. 2. Sensing a low-rank Gaussian signal of dimension $n = 100$ and with about 70% eigenvalues of Σ zero: (a) and (c) compare recovery error $\|x - \widehat{x}\|_2$ for the Info-Greedy Sensing and random sensing A, in the presence of white noise added after the measurement, and colored noise added prior to the measurement ("noise folding"), respectively; (c) and (d) show ordered number of measurements for Info-Greedy Sensing in the two cases. Info-Greedy Sensing and batch method perform identical in this case.

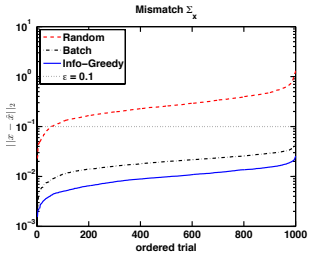

Fig. 3. Sensing a low-rank Gaussian signal of dimension $n = 500$ and about 5% of the eigenvalues are non-zero, when there is mismatch between the assumed covariance matrix and true covariance matrix: $\Sigma_{,\text{assumed}} = \Sigma_{,\text{true}} + ee^\intercal$, where $e \sim \mathcal{N}(0, I)$, and using 20 measurements. The batch method measures using the largest eigenvectors of $\Sigma_{,\text{assumed}}$, and the Info-Greedy Sensing updates $\Sigma_{,\text{assumed}}$ in the algorithm. Info-Greedy Sensing is more robust to mismatch than the batch method.

well compared with the gradient descent approach. For GMM signals, the complexity of the batch method is $\mathcal{O}(Cn^3)$ (due to eigendecomposition of C components), versus the complexity of Info-Greedy Sensing algorithm is on the order of $\mathcal{O}(Ctmn^2)$ where t is the number of iterations needed to compute the eigenvector associated with the largest eigenvector (e.g., using the power method), and m is the number of measures which is typically on the order of k.

3) Sparse Info-Greedy Sensing: Consider designing a sparse Info-Greedy Sensing vector for a single Gaussian signal with $n = 10$, desired sparsity of measurement vector $k_0 = 5$, and the low-rank covariance matrix is generated as before by thresholding eigenvalues. Fig. 8(a) shows the pattern of non-zero entries from measurement 1 to 5. Fig. 8(b) compares the performance of randomly selecting 5 non-zero entries. The sparse Info-Greedy Sensing algorithm outperforms the random

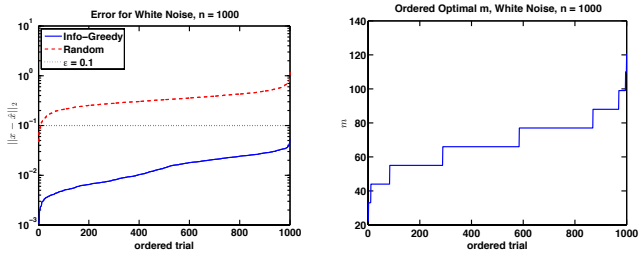

Fig. 4. Sense a low-rank Gaussian signal of dimension $n = 1000$ and about 5% eigenvalues of Σ are non-zero. Info-Greedy Sensing has two orders of magnitude improvement over the random projection.

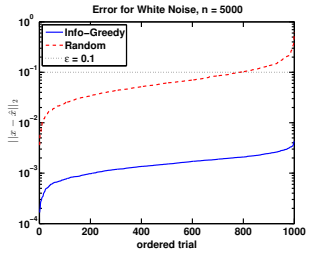

Fig. 5. Sense a Gaussian signal of dimension $n = 5000$. The covariance matrix is low-rank and sparse: only 0.0003% of entires Σ are non-zero and the rank is 3. Info-Greedy Sensing has two orders of magnitude improvement over the random projection. The number of measurements is 3 as calculated through (11).

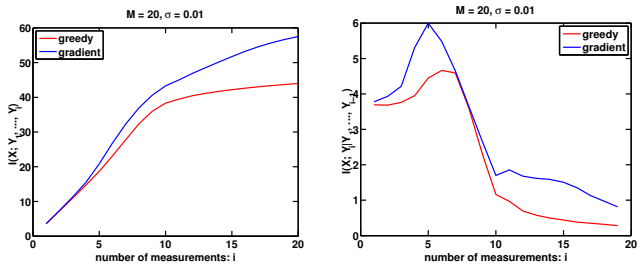

(a) (b)

Fig. 6. Sensing a GMM signal: comparison of greedy heuristic and the gradient descent approach in terms of (a) mutual information $\mathbb{I}[x; y_1, \ldots, y_i]$ over number of measurements i, average over 100 Monte Carlo trials; (b) $\mathbb{I}[x; y_i | y_1, \ldots, y_i]$ over number of measurements i, averaged over 100 Monte Carlo trials.

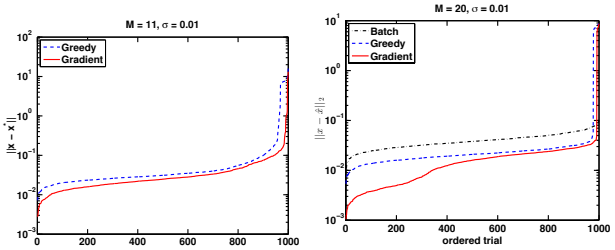

Fig. 7. Sensing a GMM signal: comparison of errors for the batch gradient descent method [32] and the Info-Greedy Sensing algorithms: the greedy heuristic and the gradient descent approach, when $m = 11$ and $m = 20$, respectively.

approach and does not degrade too much from the non-sparse Info-Greedy Sensing.

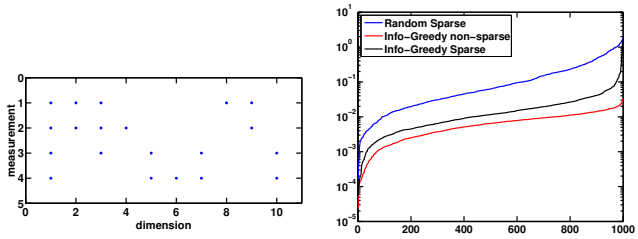

(a) (b)

Fig. 8. Results of designing sparse sensing vectors: (a) support of the sparse measurements for $n = 10$, $k_0 = 5$, over 5 measurements; (b) comparison of errors for the random sparse measurement, sparse Info-Greedy measurement, and non-sparse Info-Greedy measurement.

B. Real data

1) MNIST handwritten dataset: We exam the performance of using GMM Info-Greedy Sensing on MNIST handwritten dataset[5]. In this example, since the true label of the training data is known, we can use training data to estimate the true prior distribution π_c, μ_c and Σ_c (there are $C = 10$ classes of Gaussian components each corresponding to one digit) using 10,000 training pictures of handwritten digits picture of dimension 28 by 28. The images are vectorize and hence $n = 784$, and the digit can be recognized using the its highest posterior π_c after sequential measurements. Fig. 9 demonstrates an instance of recovered image (true label is 2) using $m = 40$ sequential measurements, for the greedy heuristic and the gradient descent approach, respectively. In this instance, the greedy heuristic classifies the image erroneously as 6, and the gradient descent approach correctly classifies the image as 2. Table I shows the probability of false classification for the testing data, where the random approach is where a_i are normalized random Gaussian vectors. Again, the greedy heuristic has good performance compared to the gradient descent method.

Fig. 9. Comparison of true and recovered handwritten digit 2 by the greedy heuristic and the gradient descent approach, respectively.

TABLE I
COMPARISON OF PROBABILITY OF FALSE CLASSIFICATION FOR MNIST HANDWRITTEN DIGITS DATASET.

Method	Random	Greedy	Gradient
prob. false classification	0.192	0.152	0.144

2) Recovery of power consumption vector: We consider recovery of a power consumption vector for 58 counties in California[6]. Data for power consumption in these counties from year 2006 to year 2012 are available. We first fit a single

[5]http://yann.lecun.com/exdb/mnist/
[6]http://www.ecdms.energy.ca.gov/elecbycounty.aspx

Gaussian model using data from year 2006 to 2011 (Fig. 10(a), the probability plot demonstrates that Gaussian is a reasonably good fit to the data), and then test the performance of the Info-Greedy Sensing in recovering the data vector of year 2012. Fig. 10(b) shows that even by using a coarse estimate of the covariance matrix from limited data (5 samples), Info-Greedy Sensing can have better performance than the random algorithm. This example has some practical implications: the compressed measurements here correspond to collecting the total power consumption over a region of the power network. This collection process can be achieved automatically by new technologies such as the wireless sensor network platform using embedded RFID in [2] and, hence, our Info-Greedy Sensing may be an efficient solution to monitoring of power consumption of each node in a large power network.

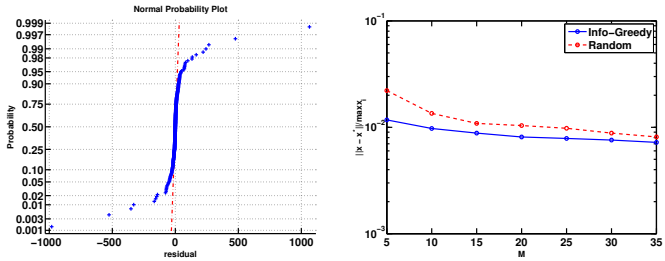

(a) (b)

Fig. 10. Recovery of power consumption data of 58 counties in California: (a) normal probability of residuals formed by training data after subtracting the mean estimated from year 2006 to year 2011; (b) relative error $\|x-\hat{x}\|_2/\|x\|_\infty$ for estimating power consumption vector in year 2012 versus the number of measurements.

VII. Conclusion

We have presented a general framework for sequential adaptive compressed sensing, Info-Greedy Sensing, which is based on maximizing mutual information between the measurement and the signal model conditioned on previous measurements. Our results demonstrate that adaptivity helps when prior distributional information of the signal is available and Info-Greedy is an efficient tool to explore such prior information, such as in the case of the GMM signals. Adaptivity also brings robustness when there is mismatch between the assumed and true distribution, and we have demonstrated such benefits for Gaussian signals. Moreover, Info-Greedy Sensing shows significant improvement over random projection for signals with sparse and low-rank covariance matrices, which demonstrate the potential value of Info-Greedy Sensing for big data.

References

[1] D. J. Brady, *Optical imaging and spectroscopy*. Wiley-OSA, April 2009.

[2] W. Boonsong and W. Ismail, "Wireless monitoring of household electrical power meter using embedded RFID with wireless sensor network platform," *Int. J. Distributed Sensor Networks, Article ID 876914, 10 pages*, vol. 2014, 2014.

[3] B. Zhang, X. Cheng, N. Zhang, Y. Cui, Y. Li, and Q. Liang, "Sparse target counting and localization in sensor networks based on compressive sensing," in *IEEE Int. Conf. Computer Communications (INFOCOM)*, pp. 2255 – 2258, 2014.

[4] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Info. Theory*, vol. 52, pp. 5406–5425, Dec. 2006.

[5] D. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1289–1306, Apr. 2006.

[6] Y. C. Eldar and G. Kutyniok, eds., *Compressed sensing: theory and applications*. Cambridge University Press Cambridge, UK, 2012.

[7] J. Haupt, R. M. Castro, and R. Nowak, "Distilled sensing: adaptive sampling for sparse detection and estimation," *IEEE Trans. Info. Theory*, vol. 57, pp. 6222–6235, Sept. 2011.

[8] D. Wei and A. O. Hero, "Multistage adaptive estimation of sparse signals," *IEEE J. Sel. Topics Sig. Proc.*, vol. 7, pp. 783 – 796, Oct. 2013.

[9] D. Wei and A. O. Hero, "Performance guarantees for adaptive estimation of sparse signals," *arXiv:1311.6360v1*, 2013.

[10] M. L. Malloy and R. Nowak, "Sequential testing for sparse recovery," *IEEE Trans. Info. Theory*, vol. 60, no. 12, pp. 7862 – 7873, 2014.

[11] E. Arias-Castro, E. J. Candès, and M. A. Davenport, "On the fundamental limits of adaptive sensing," *IEEE Trans. Info. Theory*, vol. 59, pp. 472–481, Jan. 2013.

[12] P. Indyk, E. Price, and D. P. Woodruff, "On the power of adaptivity in sparse recovery," in *IEEE Foundations of Computer Science (FOCS)*, Oct. 2011.

[13] M. L. Malloy and R. Nowak, "Near-optimal adaptive compressed sensing," *arXiv:1306.6239v1*, 2013.

[14] C. Aksoylar and V. Saligrama, "Information-theoretic bounds for adaptive sparse recovery," *arXiv:1402.5731v2*, 2014.

[15] G. Yu and G. Sapiro, "Statistical compressed sensing of Gaussian mixture models," *IEEE Trans. Sig. Proc.*, vol. 59, pp. 5842 – 5858, Dec. 2011.

[16] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. Sig. Proc.*, vol. 56, no. 6, pp. 2346–2356, 2008.

[17] J. Haupt, R. Nowak, and R. Castro, "Adaptive sensing for sparse signal recovery," in *IEEE 13th Digital Signal Processing Workshop and 5th IEEE Signal Processing Education Workshop (DSP/SPE)*, pp. 702 – 707, 2009.

[18] M. A. Davenport and E. Arias-Castro, "Compressive binary search," *arXiv:1202.0937v2*, 2012.

[19] A. Tajer and H. V. Poor, "Quick search for rare events," *arXiv:1210:2406v1*, 2012.

[20] D. Malioutov, S. Sanghavi, and A. Willsky, "Sequential compressed sensing," *IEEE J. Sel. Topics Sig. Proc.*, vol. 4, pp. 435–444, April 2010.

[21] J. Haupt, R. Baraniuk, R. Castro, and R. Nowak, "Sequentially designed compressed sensing," in *Proc. IEEE/SP Workshop on Statistical Signal Processing*, 2012.

[22] S. Jain, A. Soni, and J. Haupt, "Compressive measurement designs for estimating structured signals in structured clutter: A Bayesian experimental design approach," *arXiv:1311.5599v1*, 2013.

[23] A. Krishnamurthy, J. Sharpnack, and A. Singh, "Recovering graph-structured activations using adaptive compressive measurements," in *Annual Asilomar Conference on Signals, Systems, and Computers*, Sept. 2013.

[24] E. Tánczos and R. Castro, "Adaptive sensing for estimation of structure sparse signals," *arXiv:1311.7118*, 2013.

[25] A. Soni and J. Haupt, "On the fundamental limits of recovering tree sparse vectors from noisy linear measurements," *IEEE Trans. Info. Theory*, vol. 60, no. 1, pp. 133–149, 2014.

[26] A. Ashok, P. Baheti, and M. A. Neifeld, "Compressive imaging system design using task-specific information," *Applied Optics*, vol. 47, no. 25, pp. 4457–4471, 2008.

[27] J. Ke, A. Ashok, and M. Neifeld, "Object reconstruction from adaptive compressive measurements in feature-specific imaging," *Applied Optics*, vol. 49, no. 34, pp. 27–39, 2010.

[28] A. Ashok and M. A. Neifeld, "Compressive imaging: hybrid measurement basis design," *J. Opt. Soc. Am. A*, vol. 28, no. 6, pp. 1041– 1050, 2011.

[29] M. Seeger, H. Nickisch, R. Pohmann, and B. Schoelkopf, "Optimization of k-space trajectories for compressed sensing by Bayesian experimental design," *Magnetic Resonance in Medicine*, 2010.

[30] R. Waeber, P. Frazier, and S. G. Henderson, "Bisection search with noisy responses," *SIAM J. Control and Optimization*, vol. 51, no. 3, pp. 2261–2279, 2013.

[31] J. M. Duarte-Carvajalino, G. Yu, L. Carin, and G. Sapiro, "Task-driven adaptive statistical compressive sensing of Gaussian mixture models," *IEEE Trans. Sig. Proc.*, vol. 61, no. 3, pp. 585–600, 2013.

[32] W. Carson, M. Chen, R. Calderbank, and L. Carin, "Communication inspired projection design with application to compressive sensing," *SIAM J. Imaging Sciences*, 2012.

[33] L. Wang, D. Carlson, M. D. Rodrigues, D. Wilcox, R. Calderbank, and L. Carin, "Designed measurements for vector count data," in *Neural Information Processing Systems Foundation (NIPS)*, 2013.

[34] L. Wang, A. Razi, M. Rodrigues, R. Calderbank, and L. Carin, "Nonlinear information-theoretic compressive measurement design," in *Proc. 31st Int. Conf. Machine Learning (ICML)*, ., 2014.

[35] G. Braun, C. Guzmán, and S. Pokutta, "Unifying lower bounds on the oracle complexity of nonsmooth convex optimization," *arXiv:1407.5144*, 2014.

[36] E. Arias-Castro and Y. C. Eldar, "Noise folding in compressed sensing," *IEEE Signal Processing Letter*, vol. 18, pp. 478 – 481, June 2011.

[37] M. Chen, *Bayesian and Information-Theoretic Learning of High Dimensional Data*. PhD thesis, Duke University, 2013.

[38] Y. C. Eldar and G. Kutyniok, *Compresssed Sensing: Theory and Applications*. Cambridge Univ Press, 2012.

[39] M. Payaró and D. P. Palomar, "Hessian and concavity of mutual information, entropy, and entropy power in linear vector Gaussian channels," *IEEE Trans. Info. Theory*, pp. 3613–3628, Aug. 2009.

[40] M. A. Iwen and A. H. Tewfik, "Adaptive group testing strategies for target detection and localization in noisy environment," *Institute for Mathematics and Its Applications (IMA) Preprint Series # 2311*, 2010.

[41] Y. Xie, Y. C. Eldar, and A. Goldsmith, "Reduced-dimension multiuser detection," *IEEE Trans. Info. Theory*, vol. 59, pp. 3858 – 3874, June 2013.

[42] L. Trevisan, *Lecture notes for CS359G: Graph Partitioning and Expanders*. Stanford University, Stanford, CA, 2011.

[43] D. Palomar and S. Verdú, "Gradient of mutual information in linear vector Gaussian channels," *IEEE Trans. Info. Theory*, pp. 141–154, 2006.

[44] F. Renna, R. Calderbank, L. Carin, and M. R. D. Rodrigues, "Reconstruction of signals drawn from a Gaussian mixture via noisy compressive measurements," *IEEE Trans. Sig. Proc.*, arXiv:1307.0861. To appear.

[45] S. Arora, E. Hazan, and S. Kale, "The multiplicative weights update method: A meta-algorithm and applications," *Theory of Computing*, vol. 8, no. 1, pp. 121–164, 2012.

[46] M. A. Duran and I. E. Grossmann, "An outer-approximation algorithm for a class of mixed-integer nonlinear programs," *Math. Programming*, vol. 36, no. 3, pp. 307 – 339, 1986.

[47] A. Schrijver, *Theory of linear and integer programming*. Wiley, 1986.

[48] J. C. Duchi and M. J. Wainwright, "Distance-based and continuum Fano inequalities with applications to statistical estimation," *arXiv:1311.2669v2*, 2013.

[49] J. M. Leiva-Murillo and A. Artes-Rodriguez, "A Gaussian mixture based maximization of mutual information for supervised feature extraction," *Lecture Notes in Computer Science, Independent Component Analysis and Blind Signal Separation*, vol. 3195, pp. 271 – 278, 2004.

APPENDIX A
GENERAL PERFORMANCE LOWER BOUNDS

In the following we establish a general lower bound for the number of sequential measurements needed to obtain certain small recovery error $\|x - \widehat{x}\|_2$, similar to the approach in [35]. We consider the following model: sequentially perform measurements and performance is measured by the number M of measurements required to obtain a reconstruction of the signal with a prescribed accuracy. Assume the sequential measurements a_i are linear and the measurement returns $a_i^\intercal x$. Formally, let $\mathcal{F}$ be a finite family of signals of interest, and $F \in \mathcal{F}$ be a random variable with uniform distribution on $\mathcal{F}$. Denote by $A = (a_1, a_2, \dots)$ the sequence of measurements, and $y = (y_1, y_2, \dots)$ the sequence of measurement values: $y_i = a_i^\intercal x$. Let $\Pi = (A, y)$ denote the transcript of the measurement operations and $\Pi_i = (a_i, y_i)$ a single measurement/value pair. Note that Π is a random variable of the picked signal F. Assume that the accuracy ε is high enough to ensure a one-to-one correspondence between signal F and the ε-ball it is contained in. Thus we can return the center of such an ε-ball as the reconstruction $\widehat{x}$ of x. In this regime, an ε-recovery of a signal x is (information-theoretically) equivalent to learning the ε-ball that x is contained in, and we can invoke the *reconstruction principle*

$$\mathbb{I}[F;\Pi] = \mathbb{H}[F] = \log|\mathcal{F}|, \tag{18}$$

i.e., the transcript has to contain the same information as F and in fact uniquely identify it. With this model it was shown in [35] that the total amount of information acquired, $\mathbb{H}[F]$, is equal to the sum of the conditional information per iteration:

Theorem A.1 ([35])**.**

$$\mathbb{I}[F;\Pi] = \sum_{i=1}^{\infty} \underbrace{\mathbb{H}\left[y_i \,\middle|\, a_i, \Pi^{i-1}, M \geq i\right]}_{\text{information gain by measurement } i} \mathbb{P}[M \geq i], \tag{19}$$

where Π^{i-1} is a shorthand for $\Pi^{i-1} \triangleq (\Pi_1, \dots, \Pi_{i-1})$ and M is the random variable of required measurements.

We will use Theorem A.1 to establish Lemma III.4 that the bisection algorithm is Info-Greedy for k-sparse signals. *A priori*, Theorem A.1 does not give a bound on the expected number of required measurements, and it only characterizes how much information the sensing algorithm learns from each measurement. However, if we can upper bound the information acquired in each measurement by some constant, this leads to a lower bound on the *expected* number of measurements, as well as a *high-probability* lower bound:

Corollary A.2 (Lower bound on number of measurements)**.** *Suppose that for some constant $C > 0$,*

$$\mathbb{H}\left[y_i \,\middle|\, a_i, \Pi^{i-1}, m \geq i\right] \leq C$$

for every round i where M is as above. Then $\mathbb{E}[M] \geq \frac{\log|\mathcal{F}|}{C}$. Moreover, for all t we have $\mathbb{P}[M < t] \leq (Ct)/\mathbb{H}[F]$ and $\mathbb{P}[T = \mathcal{O}(\mathbb{H}[F])] = 1 - o(1)$.

The information theoretic approach also lends itself to lower bounds on the number of measurements for Gaussian signals, as e.g., in [48, Corollary 4].

APPENDIX B
DERIVATION OF GAUSSIAN SIGNAL MEASURED WITH COLORED NOISE

First consider the case when colored noise is added after the measurement: $y = Ax + w$, $w \sim \mathcal{N}(0, \Sigma_w)$. In the following, we assume the noise covariance matrix Σ_w is full rank. Note that we can write $w_1 = e_1^\intercal w$. Let the eigendecomposition of the noise covariance matrix be $\Sigma_w = U_w \Lambda_w U_w^\intercal$, and define a constant vector $b \triangleq \Lambda_w^{1/2} U_w^\intercal e_1$. So the variance of w_1 is given by $e_1^\intercal \Sigma_w e_1 = b^\intercal b$. Re-parameterize a_1 by introducing a unitary matrix R: $a_1 = (\sqrt{\beta_1}/\|b\|_2) Rb$. Also let the eigendecomposition of Σ be $\Sigma = U_x \Lambda_x U_x^\intercal$. Then the mutual information of x and y_1 can be written as

$$\begin{aligned}
\mathbb{I}[x; y_1] &= \frac{1}{2}\ln\left(\frac{a_1^\intercal \Sigma a_1}{e_1^\intercal \Sigma_w e_1} + 1\right) = \frac{1}{2}\ln\left(\frac{\beta_1}{\|b\|_2^2} \cdot \frac{b^\intercal R^\intercal \Sigma R b}{b^\intercal b} + 1\right) \\
&= \frac{1}{2}\ln\left(\frac{\beta_1}{\|b\|_2^2} \cdot \frac{e_1^\intercal U_w \Lambda_w^{1/2} R^\intercal U_x \Lambda_x U_x^\intercal R \Lambda_w^{1/2} U_w^\intercal e_1}{b^\intercal b} + 1\right) \\
&\leq \frac{1}{2}\ln\left(\frac{\beta_1}{\|\Lambda_w^{1/2} U_w^\intercal e_1\|_2^4} \cdot \|\Sigma\|\|\Sigma\|\right),
\end{aligned} \tag{20}$$

and the maximum is achieved when $R = U_x$. Hence, the Info-Greedy Sensing vector is

$$a_1 = \frac{\sqrt{\beta}_1}{\|\Lambda_w^{1/2} U_w^\intercal e_1\|_2} U_x \Lambda_w^{1/2} U_w^\intercal e_1. \tag{21}$$

Note that the solution (21) for a_1 has the interpretation of "mode matching", i.e., aligning of eigenspaces of the signal and the noise similar to that in [32] for the non-adaptive setting.

For the "noise folding" model with colored noise, $y = A(x + w)$, $w \sim \mathcal{N}(0, \Sigma_w)$, since the power of a_i does not affect SNR, we assume $\|a_1\|_2 = 1$. Let $d \triangleq \Lambda_w^{1/2} U_w^\intercal a_1$ and, hence, $a_1 = U_w \Lambda_w^{-1/2} d$. In this case

$$\begin{aligned}\mathbb{I}[x; y_1] &= \frac{1}{2} \ln \left(\frac{a_1^\intercal \Sigma a_1}{a_1^\intercal \Sigma_w a_1} + 1 \right) \\ &= \frac{1}{2} \ln \left(\frac{d^\intercal \Lambda_w^{-1/2} U_w^\intercal \Sigma U_w \Lambda_w^{-1/2} d}{d^\intercal d} + 1 \right) \le \frac{1}{2} \ln \left(\|\Sigma'\| + 1 \right),\end{aligned} \tag{22}$$

where $\Sigma' = \Lambda_w^{-1/2} U_w^\intercal \Sigma U_w \Lambda_w^{-1/2}$, and the maximum is achieved when d is the eigenvector for the largest eigenvector of Σ'. Equivalently a_1 is an eigenvector for the largest eigenvalue of $U_w \Lambda_w^{-1/2} \Sigma_x' \Lambda_w^{-1/2} U_w$ or the largest eigenvector for $\Sigma_w^{-1} \Sigma$. Note that in contrast to (21), in this case the "mode matching" is not possible because the noise covariance depends on the measurement vector a_i as well.

Appendix C
Derivation for GMM signals

Let $\widetilde{p}$ and $\widetilde{\mathbb{E}}$ denote the probability density function and expectation given $\{y_j : j < i\}$. Using [32], [37], [43], the gradient of mutual information with respect to a_i is given by

$$\frac{\mathbb{I}[x; y_i \,|\, y_j, j < i]}{\partial a_i} = \frac{E_i (a_i; y_j, j < i)^\intercal a_i}{\sigma^2}, \tag{23}$$

where $E_i(a_i; y_j, j < i) \in \mathbb{R}^{n \times n}$ is the MMSE matrix conditioned on measurements prior to i, which can be written as

$$E_i = \int \widetilde{p}(y) \cdot \underbrace{\int \widetilde{p}(x \mid y_i = y)(x - \widetilde{\mathbb{E}}[x \mid y_i = y])(x - \widetilde{\mathbb{E}}[x \mid y_i = y])^\intercal dx}_{g(y)} dy. \tag{24}$$

For GMM, a closed form formula for the integrand $g(y)$ can be derived. Note that the conditional distribution of x given the $\{y_j : j < i\}$ and $y_i = y$ turns out to be a GMM with updated parameters: mean $\widetilde{\mu}_c$, variance $\widetilde{\Sigma}_c$, and weight $\widetilde{\pi}_c$:

$$\widetilde{\mu}_c(y) = \mu_c + \Sigma_c D_i^\intercal (\widetilde{y}_i - D_i \mu_c)/\sigma^2, \tag{25}$$

$$\widetilde{\Sigma}_c = \Sigma_c - \Sigma_c D_i^\intercal D_i \Sigma_c / \sigma^2, \tag{26}$$

$$\widetilde{\pi}_c \propto \pi_c \Phi(\widetilde{y}_i; D_i \mu_c, D_i \Sigma_c D_i^\intercal + \sigma^2), \tag{27}$$

where $D_i^\intercal = [a_1, \cdots, a_{i-1}, a_i]$ and $\widetilde{y}_i = [y_1, \cdots, y_{i-1}, y]^\intercal$, and hence $\widetilde{\mathbb{E}}[x \mid y_i = y, c] = \widetilde{\mu}_c(y)$, $\widetilde{\mathbb{E}}[x \mid y_i = y] = \sum_{c=1}^C \widetilde{\pi}_c \widetilde{\mu}_c(y)$. Based on the above results

$$\begin{aligned} &g(y) \\ &= \sum_{c=1}^C \widetilde{\pi}_c \{ \widetilde{\Sigma}_c + (\widetilde{\mu}_c(y) - \sum_{c=1}^C \widetilde{\pi}_c \widetilde{\mu}_c(y))(\widetilde{\mu}_c(y) - \sum_{c=1}^C \widetilde{\pi}_c \widetilde{\mu}_c(y))^\intercal \}. \end{aligned} \tag{28}$$

The closed form expression (28) enables the gradient to be evaluated efficiently by drawing samples from $p(y)$ and computing direct Monte Carlo integration, as summarized in Algorithm 4. We stop the gradient descent iteration whenever the difference between two conditional mutual information drops below a threshold. The conditional mutual information for GMM is given by

$$\mathbb{I}[x; y_i \,|\, y_j, j < i] = \mathbb{H}[x \,|\, y_j, j < i] - \mathbb{H}[x \,|\, y_j, j \le i]. \tag{29}$$

Since the posterior distribution of x conditioned on $\{y_j, j < i\}$ and x conditioned on $\{y_j, j \le i\}$ are both GMM, an approximation for the entropy of GMM will approximate (29). Such an approximation is derived in [49]. For GMM described in (13) $\mathbb{H}[x] \approx \sum_{i=1}^C \pi_i \log \left((2\pi e)^{n/2} |\Sigma_i|^{1/2} / \pi_i - (C-1) \right)$. This approximation is good when the Gaussian components are not overlapping too much, or more precisely, when $\sum_{c \neq i} \pi_c \mathcal{N}(\mu_c, \Sigma_c) / (\pi_c \mathcal{N}(\mu_c, \Sigma_c)) \ll 1$.

Appendix D
Proofs

Proof of Lemma III.1: We will first prove the noiseless case. The set L is intended to consist of at most k disjoint subsets covering the part of the signal x that has not been determined yet.

At each iteration of the loop starting at Line 4, Algorithm 2 first splits every set in L into two of almost equal size, and decides which of the new sets S intersects the support of S by measuring $a_S^\intercal x$. Then keeps only the S in L, which intersect the support of x, and have size greater than 1. On the removed subsets S, the measurement $a_S^\intercal x$ already determines x, and the estimator $\widehat{x}$ is updated to coincide with x.

Now we estimate the number of measurements altogether. As the support of x has size at most k, at every iteration L consist of at most k sets, meaning $2k$ measurement per iteration. Finally, due to halving of sets, as the sizes of the sets in L are at most $2^{\lceil \log n \rceil - i}$ after iteration i, therefore after at most $\lceil \log n \rceil$ iteration, all the sets in L will have size 1, and the algorithm stops, having determined the whole x. Thus, at most $2k \lceil \log n \rceil$ measurements are made altogether.

In the noisy case, the main difference is that every measurement is repeated $r = \lceil \log n \rceil$ times, and average is taken over the block of r measurements to reduce the error to at most ε with error probability at most $\exp(-r\varepsilon^2/2\sigma^2)/2$. Assuming the error is less than ε for every block of measurements, the algorithm always correctly detects when a subset S does not intersect the support of x, as then $y \le \varepsilon$ in Line 11. On such subsets x is estimated by 0, which is exact. However, $y \le \varepsilon$ might also happen if $a_S x \le \varepsilon$ but x is not 0 on S. This will not cause L to consist of more than k subsets, but x will be estimated by 0 on S, causing errors at most ε on the non-zero coordinates on x in S. Note that Line 15 establishes an error

at most ε on subsets S with $|S| = 1$. All in all, the algorithm terminates after at most $2k\lceil \log n \rceil$ blocks of measurements, and in the end the estimator $\widehat{x}$ coincides with x outside the support of x, has error at most ε in every coordinate in the support of x. Therefore $\|\widehat{x} - x\|_2 \leq \varepsilon\sqrt{k}$. By the union bound, the probability of making an error greater than ε in some of the first $2k\lceil \log n \rceil$ blocks is at most $2k\lceil \log n \rceil \exp(-r\varepsilon/(2\sigma^2))/2$, which provides the claimed error probability. ∎

Proof of Lemma III.2: We consider the family $\mathcal{F}$ of signals consisting of all k-sparse signals on n bits with uniform distribution. In particular $\log|\mathcal{F}| = \log\binom{n}{k} > \log(\frac{n}{k})^k = k\log\frac{n}{k}$. Observe that for every measurement a we have $y = a^\intercal x \in \{0, \dots, k\}$ and hence the entropy of the measurement result is less than $\mathbb{H}[y] \leq \log k + 1$. We apply Theorem A.1 to obtain a lower bound on the expected number of measurements m:

$$\begin{aligned}\mathbb{E}[m] &\geq \frac{k\log\frac{n}{k}}{\log k + 1} = \frac{k\log n - k\log k}{\log k + 1} \\ &> \frac{k}{\log k + 1}\log n - k > \frac{k}{\log k + 1}(-1 + \log n).\end{aligned} \tag{30}$$

∎

Proof of Lemma III.3: It is easy to observe that after a measurement the size of the domain is effectively reduced, however the signal x is still distributed uniformly at random in the residual set. Thus it suffices to consider a single measurement. Let the measurement a be chosen such that the first half of the entries are 0 and the other half of the entries are 1, i.e., we partition $[n] = A_1 \dot\cup A_2$. The obtained measurement Y satisfies $Y = \begin{cases} 0, & \text{w.p. } 1/2; \\ 1, & \text{w.p. } 1/2. \end{cases}$ Note that Y is determined by X given the measurement, i.e., $\mathbb{H}[Y \mid X] = 0$. We therefore obtain $\mathbb{I}[X;Y] = \mathbb{H}[Y] - \mathbb{H}[Y \mid X] = \mathbb{H}[Y] = 1$. On the other hand $\mathbb{I}[X;Y] \leq \mathbb{H}[Y] \leq 1$ as Y is binary. Thus the measurement maximizes the mutual information. As the reduced problem after the measurement is identical to the original one except for the domain size being reduced by a factor of $1/2$, by induction, we obtain that the continued bisections maximize the conditional mutual information. ∎

Proof of Lemma III.4: The considered family $\mathcal{F}$ has entropy at least $\mathbb{H}[\mathcal{F}] = \log|\mathcal{F}| = \log\binom{n}{k} > \log(\frac{n}{k})^k = k\log\frac{n}{k}$. On the other hand, the bisection algorithm requires $k\lceil \log n \rceil$ queries. Using Theorem A.1 and let c be the upper bound on information gathered per measurement. We obtain $k\lceil \log n \rceil \geq (k\log\frac{n}{k})/c$. Solving for c we obtain $c \geq 1 - (\log k)/(\log n)$. Thus the expected amount of information per query is at least c. ∎

Proof of Theorem IV.1: We consider how the covariance matrix of x changes conditioned on the measurements taken. As explained in (10), measuring with an eigenvector reduces its eigenvalue from λ to $\lambda\sigma^2/(\lambda + \sigma^2)$ leaving the other eigenvalues unchanged. Thus, as far as the spectrum of the covariance matrix is concerned, each measurements applies this reduction to one of the then-largest eigenvalue. Note that Algorithm 3 might reduce several times an eigenvalue, but as mentioned before, several reductions has the same effect as one reduction with the combined power. Thus, to reduce λ_i to a value at most $\delta \coloneqq \varepsilon^2/\chi^2_n(p)$, the minimum required number of measurements is $(1/\delta - 1/\lambda_i)\,\sigma^2$ provided $\lambda_i > \delta$ and $\sigma > 0$. Rounding up to integer values and summing up for all directions, we obtain (9a) as a lower bound on total power. Furthermore, if $\sigma^2 \leq \delta$ then a single measurement suffices to ensure $\lambda_i \leq \delta$, and a measurement is only needed if $\lambda_i > \delta$. This provides (9b) even in the noiseless case $\sigma = 0$. All in all, after the algorithm has finished, the posterior distribution of the signal x is Gaussian $\mathcal{N}(\mu', \Sigma')$ with mean μ' and covariance matrix Σ'. The largest eigenvalue $\|\Sigma'\|$ of Σ' is at most δ, i.e., $\varepsilon \geq \sqrt{\|\Sigma'\| \cdot \chi^2_n(p)}$. As a consequence, we show that the mean μ' returned by the algorithm is an estimator of the signal with the required accuracy ε. An easy calculation shows that the distance between x and μ' is at most ε with probability at least p:

$$\begin{aligned}&\mathbb{P}_{x\sim\mathcal{N}(\mu',\Sigma')}\left[\|x - \mu'\|_2 \leq \varepsilon\right] \\ &\geq \mathbb{P}_{x\sim\mathcal{N}(\mu',\Sigma')}\left[\|x - \mu'\|_2 \leq \sqrt{\|\Sigma'\| \cdot \chi^2_n(p)}\right] \\ &\geq \mathbb{P}_{x\sim\mathcal{N}(\mu',\Sigma')}\left[(x-\mu')^\intercal \Sigma'^{-1}(x - \mu') \leq \chi^2_n(p)\right] = p,\end{aligned}$$

where the last equality is a restatement of the well-known prediction interval for multivariate normal distributions. ∎

Proof of Theorem IV.2: The proof is similar to that of Theorem IV.1, so we point out only the differences. The power used by Algorithm 3 reduced every eigenvalue λ_i to exactly $\delta \coloneqq \frac{\varepsilon^2}{\chi^2_n(p)}$, provided $\lambda_i > \delta$, otherwise λ_i is left intact. Hence summing up the powers for the eigenvalues, the total is power is given by (11), and the largest eigenvalue of the posterior covariance matrix is at most δ. The mean is an estimator of the signal with the required accuracy for the same reasons as in the proof of Theorem IV.1. ∎

Proof of Theorem IV.3: The proof is similar to Theorem IV.1. The only difference is that instead of the canonical scalar product the one with matrix Σ_w is used. To make this transparent, we switch to an orthonormal basis of Σ_w, we write $\Sigma_w = F^\intercal F$, and use F as a change of basis: thus the signal in the new basis is $F^{-1}x$, the measurement vectors are $F^\intercal a_i$, the covariance matrices are $\Sigma_w^{-1}\Sigma$ for the signal $F^{-1}x$, and the identity matrix for the noise. In this basis, the algorithm is identical to that of for the white noise added prior to measurement case in Algorithm 3, and hence reduces every eigenvalue of $\Sigma_w^{-1}\Sigma$ to be at most $\|\Sigma_w\|^{-1}\frac{\varepsilon^2}{\chi^2_n(p)}$. Note that the noise model is $y_i = (F^\intercal a_i)^\intercal(F^{-1}x) + (F^\intercal\widehat{a}_i)^\intercal(F^{-1}w)$, and therefore the power β_i provided by the formula $F^\intercal a_i = \sqrt{\beta_i}F^\intercal\widehat{a}_i$, i.e., $a_i = \sqrt{\beta_i}\widehat{a}_i$. In other words, the power β_i is still the length of a_i in the original basis. Let Σ' denote the posterior covariance matrix of x in the original basis. Hence $\|\Sigma_w^{-1}\Sigma'\| \leq \|\Sigma_w\|^{-1}(\varepsilon^2/\chi^2_n(p))$, and therefore $\|\Sigma'\|$ is at most $(\varepsilon^2/\chi^2_n(p))$. This ensures that the posterior mean of x returned by the algorithm to be of the required accuracy, as in Theorem IV.1. ∎

Sketch of proof for Theorem IV.5: We first sequentially run the corresponding algorithm for each component $c \in C$. This leads to a number of iteration (or power consumption) of at most $\sum_{c\in C} m_c$. We perform measurements that maximize the mutual information between the signal and the measurement outcome for the mixture. With each measurement a_i and the outcome y_i, we update the posterior distribution of the Gaussian

component, which we index by π_c^i, $c = 1, \ldots, C$, as follows

$$\pi_c^{i+1} \coloneqq \pi_c^i \cdot K_i \cdot e^{-\frac{1}{2}\frac{(y_i - a_i^\intercal \mu_{c,i})^2}{a_i^\intercal \Sigma_{c,i} a_i + \sigma^2}}, \tag{31}$$

where $\mu_{c,i}$ is the posterior mean of component c obtained after the measurement, and K_i is a normalization ensuring that π^{i+1} sum up to 1. The updates in (31) scale down the probabilities of those components c whose mean $\mu_{c,i}$ leads to reconstruction with a higher error, which is measured by $\frac{(y_i - a_i^\intercal \mu_{c,i})^2}{a_i^\intercal \Sigma_{c,i} a_i + \sigma^2}$. Then we apply the hedge version of the multiplicative weight update formula (see e.g., [45, Theorem 2.3]) to our setup and obtain that after m measurements we have

$$\frac{1}{m}\left|\sum_{i=1}^{m} K_i \cdot \left[\left(\sum_{\ell=1}^{C} \frac{(y_i - a_i^\intercal \mu_{\ell,i})^2}{a_i^\intercal \Sigma_{\ell,i} a_i + \sigma^2} \cdot \pi_\ell^i\right) - \frac{(y_i - a_i^\intercal \mu_{c,i})^2}{a_i^\intercal \Sigma_{c,i} a_i + \sigma^2}\right]\right| \leq \tilde{\eta} + \frac{2 \ln |C|}{m},$$

for all $c \in C$. Here $\tilde{\eta} > 0$ is a parameter for multiplicative update algorithm. In particular we can identify the correct component c^* whenever $m = \mathcal{O}(\frac{1}{\tilde{\eta}} \ln |C|)$. ■